Bayesian Model-Averaged Benchmark Dose Analysis Via Reparameterized Quantal-Response Models

Qijun Fang[a], **Walter W. Piegorsch**[a,b], **Susan J. Simmons**[c], **Xiaosong Li**[c], **Cuixian Chen**[c] and **Yishi Wang**[c]

[a]*Graduate Interdisciplinary Program in Statistics* and [b]*BIO5 Institute*
University of Arizona, Tucson, AZ, USA
[c]*Department of Mathematics and Statistics*
University of North Carolina at Wilmington, Wilmington, NC, USA

arXiv:1402.3896v1 [stat.ME] 17 Feb 2014

Abstract

An important objective in biomedical risk assessment is estimation of minimum exposure levels that induce a pre-specified adverse response in a target population. The exposure/dose points in such settings are known as Benchmark Doses (BMDs). Recently, parametric Bayesian estimation for finding BMDs has become popular. A large variety of candidate dose-response models is available for applying these methods, however, leading to questions of model adequacy and uncertainty. Here we enhance the Bayesian estimation technique for BMD analysis by applying Bayesian model averaging to produce point estimates and (lower) credible bounds. We include reparameterizations of traditional dose-response models that allow for more-focused use of elicited prior information when building the Bayesian hierarchy. Performance of the method is evaluated via a short simulation study. An example from carcinogenicity testing illustrates the calculations.

Keywords: Bayesian BMDL, Bayesian model averaging, benchmark analysis, dose-response analysis, hierarchical modeling, model uncertainty, multimodel inference, prior elicitation, quantitative risk assessment.

1 Introduction

1.1 Benchmark Risk Analysis

An important objective in quantitative risk assessment is characterization of detrimental or adverse responses after exposure to biological, chemical, physical, environmental, or other hazardous agents (Stern, 2008). In this context, the risk is often quantified via a dose-response function, $R(d)$ (also called the *risk function*), which is defined as the probability of exhibiting the adverse effect in a subject, object, or system exposed to a particular dose or exposure level, d, of the agent. For biomedical and environmental risk applications, statistical dose-response models are fit to data from bioassays on small mammals or other biological systems, or from epidemiological analyses of human populations at risk. The observations in such risk assessments are often in the form of proportions. This is the *quantal response setting*, and is our focus here.

A modern approach to quantal-response risk estimation is known as *benchmark analysis*. First introduced for toxicological applications by Crump (1984), this method uses a specific functional assumption on $R(d)$ to provide low-dose estimates for the risk. When applied in biomedical or other public health settings, however, adjustments for spontaneous effects must be introduced, re-expressing the risk in terms of an excess-above-background, exposure-related rate (Piegorsch and Bailer, 2005, §4.2.1). For instance, with quantal data the *extra risk function* $R_E(d) = \frac{R(d)-R(0)}{1-R(0)}$, where $R(0)$ is the background risk, is frequently employed. From this, the *benchmark dose* (BMD) is calculated by inverting $R_E(d)$ at a predetermined level of risk; the latter is called the *benchmark risk* or *benchmark response* (BMR). In effect, we solve for $d = \text{BMD}$ in $R_E(\text{BMD}) = \text{BMR} \in (0, 1)$; BMRs between 0.01 and 0.10 are most often seen in practice (U.S. EPA, 2012). If the exposure is measured as a concentration, the benchmark point is referred to as a *benchmark concentration* (BMC), if exposure is some sort of quantitative index the benchmark point is a *benchmark index* (BMI), etc. To illustrate, consider the following example.

Example 1. Benchmarking mammalian carcinogenicity of cumene.
Cumene, the colloquial name for isopropylbenzene (C_9H_{12}), is a hydrocarbon solvent employed in the production of industrial compounds such as phenol and acetone. Occupational and industrial exposures to cumene are common, so the U.S. National Toxicology Program (NTP) explored various forms of mammalian toxicity to the chemical (U.S. NTP, 2009). For example, Table 1 displays quantal-response data on induction of lung tumors (alveolar/bronchiolar adenomas and carcinomas) by cumene in laboratory mice after two year chronic inhalation exposure.

Table 1: Quantal carcinogenicity data: Alveolar/bronchiolar adenomas and carcinomas in female B6C3F$_1$ mice after inhalation exposure to cumene (C_9H_{12}). Source: U.S. NTP (2009).

Exposure conc. (ppm), d_i	0	125	250	500
Animals with tumors, Y_i	4	31	42	46
Animals tested, N_i	50	50	50	50

In the table, a clear dose response is evidenced. Of additional interest, however, is calculation of a benchmark exposure level to inform risk characterization of this potential carcinogen. Notice that the C_9H_{12} exposure dose, d, is actually a concentration (in ppm) here, and so technically we will compute benchmark concentrations (BMCs) based on the quantal carcinogenicity data. A benchmark analysis of these data appears Section 4.

One critical enhancement in benchmark analysis is construction of statistical confidence (or credible) intervals for the BMD to account for variability in the estimation process. Driven by public health or other safety considerations, only one-sided, lower limits are employed, denoted as BMDLs (Crump, 1995). Where needed for clarity, we add a subscript for the BMR level at which each quantity is calculated: $\text{BMD}_{100\text{BMR}}$ and $\text{BMDL}_{100\text{BMR}}$. In this fashion, BMDs and BMDLs are employed for risk characterization by a number of government and private entities. Their use for quantifying and managing risk with a variety of biological endpoints is growing in both the United States and the European Union (U.S. General Accounting Office, 2001; European Union, 2003; OECD, 2006, 2008).

1.2 Parametric Bayesian Benchmark Analysis

Statistical estimation of the BMD is fairly well-developed under a frequentist schema (Piegorsch and Bailer, 2005, §4.3); however, Bayesian benchmark analysis has only recently garnered appreciative attention. Some modern advances can be found in Naufal *et al.* (2009); Shao and Small (2011); Shao (2012); Wheeler and Bailer (2012); and Guha *et al.* (2013). Notably, these sources for calculating Bayesian BMDs generally parameterize the model in terms of standard regression-type quantities. For example, the logistic dose response $R(d) = (1+\exp\{-\beta_0-\beta_1 d\})^{-1}$ is popular, with β_0 and β_1 representing the intercept and slope, respectively, of the risk on a logit scale. Under this sort of traditional parameterization, the Bayesian hierarchy is usually presented with objective and/or improper prior distributions for the unknown β-parameters. [Some noteworthy exceptions include incorporation at the prior level of historical control information related to β_0; see Shao (2012) and Wheeler and Bailer (2012).] Indeed, it is unusual for truly informative, risk-analytic, prior knowledge to be available on regression-type parameters such as β_0 and β_1, since their interpretation is so generic. Unfortunately, this strategy can neglect informative prior information available on the true quantity of interest in this setting, the BMD.

Of course, prior information on the BMD may not always exist in practice, and in this case objective priors serve a useful purpose. When informative prior knowledge is available, however, the Bayesian paradigm can achieve its full potential. Along these lines, we previously developed reparameterized dose-response models that explicitly incorporate the BMD and other pertinent parameters into the model hierarchy (Fang and Piegorsch, 2014). When applied to Bayesian dose-response modeling, these reparameterizations allow for more practical elicitation of prior information in a benchmark analysis. Our method assumes, however, that the choice of dose-response function is made without any uncertainty, i.e., that the specification for $R(d)$—reparameterized or otherwise— is unambiguous and correct. In practice, the extensive library of quantal dose-response forms for $R(d)$ available to the risk assessor can lead to uncertainty in the model specification. To mitigate concerns over model adequacy/uncertainty, recent works employ model averaging techniques. Bayesian model

averaging (BMA; see Hoeting *et al.*, 1999) has become popular in this regard, and a number of articles have applied some form of BMA to benchmark analysis; see, e.g., Bailer *et al.* (2005); Morales *et al.* (2006); Shao and Small (2011); Shao (2012); Shao and Gift (2014); and the references therein. Most of these employ traditional β-parameterizations with objective priors, however. Here, we study the BMA approach for benchmark analysis by using the reparameterized framework from Fang and Piegorsch (2014). Section 2 reviews this reparameterized paradigm for hierarchical Bayesian benchmark analysis. Section 3 follows by connecting it with the elements of Bayesian model averaging. Section 4 returns to the cumene carcinogenicity data from Table 1 to explore practical implementation of this Bayesian model-averaged benchmark analysis, while Section 5 examines the characteristics of our BMA Bayesian BMDLs via a short simulation study. Section 6 ends with a brief discussion.

2 Hierarchical Bayesian Benchmark Analysis

2.1 Bayesian Dose-response Models

Under the quantal-response formulation, denote Y_i as the number of responses, N_i as the number of subjects tested, and $R(d_i)$ as the unknown probability that an individual subject will respond at (ordered) dose $d_i \geq 0$, $i = 1, \ldots, m$. Assume that, independently, $Y_i \sim \text{Bin}(N_i, R(d_i))$. Let $R(d_i)$ depend upon $\boldsymbol{\theta}$, an unknown parameter vector. The joint p.m.f. of $\boldsymbol{Y} = (Y_1, \ldots, Y_m)$ is then:

$$f(\boldsymbol{Y}|\boldsymbol{\theta}) = \prod_i^m \binom{N_i}{Y_i} R(d_i)^{Y_i}(1 - R(d_i))^{N_i - Y_i}. \tag{2.1}$$

Previous parametric representations for modeling $R(d)$ in risk-analytic carcinogenicity testing have focused on a suite of eight different functions (Wheeler and Bailer, 2009; Piegorsch *et al.*, 2013). These correspond to popular choices in the U.S. EPA's BMDS software (Davis *et al.*, 2012), and are illustrated in Table 2. [Throughout, $\Phi(\cdot)$ represents the c.d.f. from the standard normal distribution.] Notice in the table that models M_1–M_4 employ only two unknown parameters, while models M_5–M_8 employ three. The BMD for each model is obtained by first setting the extra risk function (as defined in Section 1) equal to the BMR and then solving for d.

In applications of Bayesian benchmark analysis to quantal data, objective forms for the prior p.d.f. $\pi(\boldsymbol{\theta})$ are common, as indicated earlier. These typically appear as diffuse Gaussian priors on the β-parameters. From this, the joint posterior distribution for $\boldsymbol{\theta}$ is obtained from Bayes formula (Casella and Berger, 2002, §7.2.3). An advantage here is that objective priors are usually easy to apply: although they generally lead to intractable integrals, computer-intensive operations such as Markov chain Monte Carlo (McMC) methods can produce a corresponding sample from the joint posterior of $\boldsymbol{\theta}$ (Robert and Casella, 2011). If the sample is sufficiently large and stable, the output can be used to approximate the posterior, from which inferences on the BMD may be conducted.

As we suggest above, a disadvantage with the traditional parameterizations in Table 2 is that the β-parameters may have unclear subject-matter interpretations if those parameters

Table 2: Common quantal-response models in toxicological and carcinogenic risk assessment. Note: BMR$\in (0,1)$ is the benchmark response and BMD is the benchmark dose

Code	Name	Risk Function, $R(d)$	BMD	Constraints
M_1	Logistic	$(1+\exp\{-\beta_0-\beta_1 d\})^{-1}$	$\frac{1}{\beta_1}\ln\left(\frac{1+\text{BMR}e^{-\beta_0}}{1-\text{BMR}}\right)$	—
M_2	Probit	$\Phi(\beta_0+\beta_1 d)$	$\frac{\Phi^{-1}\{[1-\Phi(\beta_0)]\text{BMR}+\Phi(\beta_0)\}-\beta_0}{\beta_1}$	—
M_3	Quantal-Linear	$1-\exp\{-\beta_0-\beta_1 d\}$	$-\frac{1}{\beta_1}\ln(1-\text{BMR})$	$\beta_0\geq 0,\ \beta_1\geq 0$
M_4	Quantal-Quadratic	$\gamma_0+(1-\gamma_0)(1-e^{-\beta_1 d^2})$	$\sqrt{-\frac{1}{\beta_1}\ln(1-\text{BMR})}$	$0\leq\gamma_0<1,\ \beta_1\geq 0$
M_5	Two-Stage	$1-\exp\{-\beta_0-\beta_1 d-\beta_2 d^2\}$	$\frac{-\beta_1+\sqrt{\beta_1^2-4\beta_2\ln(1-\text{BMR})}}{2\beta_2}$	$\beta_0,\beta_1,\beta_2\geq 0$
M_6	Log-Logistic	$\gamma_0+\frac{1-\gamma_0}{1+\exp(-\beta_0-\beta_1\ln\{d\})}$	$\exp\left\{\frac{\ln\left(\frac{\text{BMR}}{1-\text{BMR}}\right)-\beta_0}{\beta_1}\right\}$	$0\leq\gamma_0<1,\ \beta_1\geq 0$
M_7	Log-Probit	$\gamma_0+(1-\gamma_0)\Phi(\beta_0+\beta_1\ln\{d\})$	$\exp\left\{\frac{\Phi^{-1}(\text{BMR})-\beta_0}{\beta_1}\right\}$	$0\leq\gamma_0<1,\ \beta_1\geq 0$
M_8	Weibull	$\gamma_0+(1-\gamma_0)(1-\exp\{-e^{\beta_0}d^{\beta_1}\})$	$\exp\left\{\frac{\ln(-\ln[1-\text{BMR}])-\beta_0}{\beta_1}\right\}$	$0\leq\gamma_0<1,\ \beta_1\geq 1$

are not the target quantities of interest. If informative prior information were available on the risk-analytic quantities under study, the ambiguous interpretation(s) of these traditional, regression-type parameterizations makes incorporation of such information more difficult. This may hinder effective application of the Bayesian approach in this benchmark setting.

2.2 Reparameterizing the Quantal Response Models

For benchmark risk analysis, we argue (Fang and Piegorsch, 2014) that substantive prior knowledge may be available more often than usually considered, but not in the form of information on regression-type β-parameters. Instead, a risk assessor, toxicologist, or other domain expert would typically have prior knowledge about the target parameter, the BMD, and possibly also about other application-specific values. Our goal is to utilize the domain expert's prior knowledge for making inferences on the BMD. To do so, we follow Fang and Piegorsch (2014) and reparameterize $R(d)$ in terms of ostensibly meaningful parameters whose prior distributions are more intuitive to elicit in practice.

For the dose-response models with two parameters in Table 2, we reparameterize in terms of the target value, BMD (denoted in the sequel as ξ), and the background risk, say, $\gamma_0 = R(0)$. Thus $\boldsymbol{\theta}$ becomes the vector $[\xi\ \gamma_0]^{\mathrm{T}}$.

For the models with three unknown parameters in Table 2, we reparameterize with ξ = BMD, $\gamma_0 = R(0)$, and a parameter γ_1 defined as $R(d_\ell)$ for some non-zero dose level d_ℓ. Unless otherwise specified, we set d_ℓ to the highest dose, so $\gamma_1 = R(d_m)$. Thus, we take $\boldsymbol{\theta} = [\xi\ \gamma_0\ \gamma_1]^{\mathrm{T}}$. The latter two quantities are technically nuisance parameters as far as the BMD is concerned, but one or both are nonetheless likely to be associated with non-trivial

prior information; e.g., historical control data may inform $\gamma_0 = R(0)$ (Wheeler and Bailer, 2012; Shao, 2012).

For instance, consider a highly popular three-parameter model from carcinogenicity assessment, the two-stage version of the multi-stage model (Armitage and Doll, 1954; Nitcheva *et al.*, 2007). This is model M_5 in Table 2:

$$R(d) = 1 - \exp\{-\beta_0 - \beta_1 d - \beta_2 d^2\}, \tag{2.2}$$

where $\beta_j \geq 0$ for all $j = 0, 1, 2$. It is fairly straightforward to apply our reparameterization strategy and represent the three β-parameters in terms of $\boldsymbol{\theta} = [\xi\ \gamma_0\ \gamma_1]^{\mathrm{T}}$; we explicate this in a Supplementary Document. The result is

$$R(d) = \gamma_0 + (1 - \gamma_0)\left[1 - \exp\left\{\frac{C_5 d_m d(d_m - d) + \Gamma_5 \xi d(\xi - d)}{\xi d_m(\xi - d_m)}\right\}\right] \tag{2.3}$$

where $\Gamma_5 = \log\left(\frac{1-\gamma_1}{1-\gamma_0}\right)$ and $C_5 = -\log(1-\mathrm{BMR})$. Admittedly, this is much less compact than the common form in (2.2), but it nonetheless gives an expression that can be manipulated effectively for Bayesian benchmark analysis.

In similar fashion, the remaining models in Table 2 can be reparameterized as follows:

$\mathbf{M}_1$: $R(d) = \left[1 + \exp\left\{-\mathrm{logit}(\gamma_0) - \frac{d}{\xi}\log\left(\frac{1+\exp\{-\mathrm{logit}(\gamma_0)\}\cdot\mathrm{BMR}}{1-\mathrm{BMR}}\right)\right\}\right]^{-1}$.

$\mathbf{M}_2$: $R(d) = \Phi\left(\Phi^{-1}\{\gamma_0\} + \frac{\Phi^{-1}\{\mathrm{BMR}[1-\gamma_0]+\gamma_0\} - \Phi^{-1}(\gamma_0)}{\xi} d\right)$.

$\mathbf{M}_3$: $R(d) = 1 - \exp\left\{\log(1-\gamma_0) + \frac{\log(1-\mathrm{BMR})}{\xi} d\right\}$.

$\mathbf{M}_4$: $R(d) = \gamma_0 + (1-\gamma_0)\left(1 - \exp\left\{\frac{\log(1-\mathrm{BMR})}{\xi^2} d^2\right\}\right)$.

$\mathbf{M}_5$: See Equation (2.3).

$\mathbf{M}_6$: $R(d) = \gamma_0 + (1-\gamma_0)\left[1 + \exp\left\{\frac{C_6[\log d_m - \log d] + \Gamma_6[\log\xi - \log d]}{\log\xi - \log d_m}\right\}\right]$,
where $\Gamma_6 = \log\left(\frac{1-\gamma_1}{1-\gamma_0}\right)$ and $C_6 = \log\left(\frac{\mathrm{BMR}}{1-\mathrm{BMR}}\right)$.

$\mathbf{M}_7$: $R(d) = \gamma_0 + (1-\gamma_0)\Phi\left(\frac{C_7[\log d_m - \log d] + \Gamma_7[\log d - \log\xi]}{\log d_m - \log\xi}\right)$,
where $\Gamma_7 = \Phi^{-1}\left(\frac{\gamma_1-\gamma_0}{1-\gamma_0}\right)$ and $C_7 = \Phi^{-1}(\mathrm{BMR})$.

$\mathbf{M}_8$: $R(d) = \gamma_0 + (1-\gamma_0)\left(1 - \exp\left\{-\exp\left(\frac{C_8[\log d_m - \log d] + \Gamma_8[\log d - \log\xi]}{\log d_m - \log\xi}\right)\right\}\right)$,
where $\Gamma_8 = \log\left(-\log\left(\frac{1-\gamma_1}{1-\gamma_0}\right)\right)$ and $C_8 = \log(-\log(1-\mathrm{BMR}))$.

Here again, these reparameterizations present more burdensome notation for $R(d)$. The explicit incorporation of the target parameter ξ and well-understood quantities such as γ_0 and γ_1 allows us, however, to formulate a clearer, more application-oriented hierarchical model, from which to produce inferences on ξ.

2.3 Bayesian Benchmark Analysis under Reparameterized Quantal Response Models

Generically, under our reformulation we assign a joint p.d.f. to $\boldsymbol{\theta}$: $\pi(\boldsymbol{\theta}) = \pi(\xi, \gamma_0)$ for the two-parameter quantal-response models, or $\pi(\boldsymbol{\theta}) = \pi(\xi, \gamma_0, \gamma_1)$ for the three-parameter models. Mimicking previous Bayesian constructions for benchmark analysis (Shao and Small, 2011; Shao, 2012), we assume the unknown parameters enter into the prior independently, so that $\pi(\xi, \gamma_0) = \pi(\xi)\pi(\gamma_0)$ or $\pi(\xi, \gamma_0, \gamma_1) = \pi(\xi)\pi(\gamma_0)\pi(\gamma_1)$.

For the non-negative, target quantity ξ we employ an inverse gamma prior: $\xi \sim IG(\alpha, \beta)$ with marginal prior density $\pi(\xi|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)}\xi^{-(\alpha+1)}e^{-\beta/\xi}I_{(0,\infty)}(\xi)$, where $\Gamma(a)$ is the usual Gamma function and $I_{\mathbb{A}}(x)$ is the indicator function that returns 1 if $x \in \mathbb{A}$ and 0 otherwise. For the probability parameter $\gamma_0 = R(0)$ we take $\gamma_0 \sim Beta(\psi, \omega)$ with marginal prior $\pi(\gamma_0|\psi, \omega) = \frac{\Gamma(\psi+\omega)}{\Gamma(\psi)\Gamma(\omega)}\gamma_0^{\psi-1}(1-\gamma_0)^{\omega-1}I_{(0,1)}(\gamma_0)$. Where needed (models M_5–M_8) we similarly set $\gamma_1 \sim Beta(\kappa, \lambda)$.

The various hyperparameters—α, β, ψ, ω, κ and λ—require complete specification for implementation of the model as we propose it. To do so, we first attempt to elicit each prior by incorporating subject-matter knowledge of the associated quantities from either an individual domain expert or previous information in the literature (or both). Since the IG and beta priors each have two parameters, two associated quantities are needed. Based on interactions with toxicologists and risk assessors, we find that for the target parameter ξ, specification of the first/lower quartile (Q_1) along with the median (Q_2) of the IG prior is most propitious (Fang and Piegorsch, 2014) . Similarly, for the Beta prior on γ_0 (and γ_1) we also elicit the two quartiles Q_1 and Q_2. We then numerically solve the resulting system of equations for the pertinent hyper-parameters given the two quartiles. Details are provided in the Supplementary Document.

When elicitation is not possible, we default to objective specifications for the prior densities. We favor a simple approach: for an objective prior on ξ, use the proper prior, $\xi \sim IG(0.001, 0.001)$. This IG prior is a popular suggestion in the literature for right-skewed, positive values (Lambert *et al.*, 2005; Christensen *et al.*, 2011, §1.2), such as the BMD. It is in effect an approximation of the conventional improper reciprocal prior for positive quantities, i.e., $\pi(\xi) \propto 1/\xi$ (O'Hagan, 1994, §9.17). To ensure computational stability, following guidance from Fang and Piegorsch (2014), we perform all our hierarchical calculations with doses scaled so that the maximum administered dose equals 1. For an objective prior on either γ_0 and/or γ_1, we choose another conventional, objective prior for proportions, the univariate Jeffreys prior: $Beta\left(\frac{1}{2}, \frac{1}{2}\right)$.

From these, the posterior p.d.f. for $\boldsymbol{\theta}$ under the two-parameter dose-response models is

$$\pi(\xi, \gamma_0|\boldsymbol{Y}) = \frac{\prod_{i=1}^n \binom{N_i}{Y_i} R(d_i)^{Y_i}[1-R(d_i)]^{N_i-Y_i}}{m(\boldsymbol{Y})} \frac{\beta^\alpha e^{-\beta/\xi}}{\Gamma(\alpha)\xi^{\alpha+1}} \frac{\Gamma(\psi+\omega)}{\Gamma(\psi)\Gamma(\omega)}\gamma_0^{\psi-1}(1-\gamma_0)^{\omega-1} \quad (2.4)$$

over $\xi > 0$ and $0 < \gamma_0 < 1$. A similar form emerges when considering the three-parameter dose-response models.

The denominator of $\pi(\xi, \gamma_0|\boldsymbol{Y})$ in (2.4) contains the marginal likelihood

$$m(\boldsymbol{Y}) = \int_0^\infty \int_0^1 f(\boldsymbol{Y}|\xi, \gamma_0)\pi(\xi|\alpha, \beta)\pi(\gamma_0|\psi, \omega)d\gamma_0 d\xi,$$

where $f(\boldsymbol{Y}|\xi, \gamma_0)$ is the binomial likelihood. This marginal is intractable under our elicited prior structure, unfortunately, and to evaluate (2.4) we turn to Monte Carlo posterior approximations that produce simulated realizations $\boldsymbol{\theta}_k$, $k = 1, \ldots, K$ of the parameter vector from the posterior distribution. Our choice for the Monte Carlo technique employs an adaptive Metropolis (AM) strategy, which tunes the variance of the underlying Metropolis proposal density adaptively when generating ongoing draws of the chain. From our experience with a variety of such methods, we favor a global AM procedure with componentwise adaptive scaling described by Andrieu and Thoms (2008). Details are similar to those in Fang and Piegorsch (2014) and are given in the Supplementary Document.

We monitor convergence of the AM sample to the posterior joint distribution via standard MC diagnostics. We include a burn-in over the first $K_0 - 1$ draws, where we choose K_0 to be much smaller than K. K_0 is determined from a series of diagnostics given by Geweke (1992). (Again, details can be found in the Supplementary Document.) If the diagnostics indicate that convergence is not evidenced, we flag the result as an 'algorithm failure.' In Fang and Piegorsch (2014) we found such failures to be rare events; they usually appear with very shallow dose-response patterns. Our experience also suggests that $K = 100,000$ iterations of the AM chain, including burn-in, generally provide stable results. From these, we use the remaining $K^* = K - K_0 + 1$ draws to approximate the joint posterior for $\boldsymbol{\theta}$.

For estimating the BMD, one can follow a standard approach and select the Bayes estimator as the posterior mean of the AM sample (Casella and Berger, 2002, §7.2.3). Other formulations are also possible; see Fang and Piegorsch (2014). For the corresponding Bayesian BMDL, say, $\underline{\xi}_{100\text{BMR}}$, we find the one-sided, lower, $100(1-\alpha)\%$ credible limit on ξ, satisfying $P(\xi > \underline{\xi}_{100\text{BMR}}|\boldsymbol{Y}) = 1 - \alpha$. At the traditional level of $\alpha = 0.05$, we desire the lower 5th percentile of $\pi(\xi|\boldsymbol{Y})$, and we approximate this using the lower 5th percentile from our Monte Carlo sample of ξ. If $\{\xi_{(k)}\}_{k=1}^{K^*}$ denotes the ordered values from the retained chain, our Bayesian BMDL takes the form $\underline{\xi}_{100\text{BMR}} = \xi_{(\lfloor 0.05K^* \rfloor)}$, where $\lfloor x \rfloor$ is the floor function that returns the largest integer smaller than x.

3 Bayesian Model Averaging

The selection of parametric forms for $R(d)$ presented in §2.2 illustrates the wide variety of dose-response functions available to the risk analyst. Many of these operate well at (higher) doses near the range of the observed quantal outcomes; however, they can produce wildly different estimates on BMDs at very small levels of risk (Faustman and Bartell, 1997; Kang *et al.*, 2000). The corresponding issue of model uncertainty has bedevilled benchmark analysts since its introduction in the mid-1980s. Some users have turned to formal model selection procedures to derive the BMD (Davis *et al.*, 2012). This is a natural option, although it relies on a reliable selection criterion. In fact, model selection based on the popular Akaike information criteria (AIC) (Akaike, 1973) has been shown to select incorrect models for BMD estimation almost as often as it selects correct models (West *et al.*, 2012). As an alternative to simple model selection methods, and to provide a more model-robust option for BMD estimation, we consider here a (Bayesian) model averaging approach.

Suppose $Q > 1$ quantal-response models are under consideration such as the $Q = 8$ models in Table 2. These form an uncertainty class U_Q, with individual model elements

M_q ($q = 1, \ldots, Q$). Following Hoeting *et al.* (1999), the model-averaged posterior density for ξ is defined as a mixture of marginal posterior densities for ξ under each model, $f(\xi|\boldsymbol{Y}, M_q)$, with the individual-model posterior probabilities $P(M_q|\boldsymbol{Y})$ employed as weights:

$$f(\xi|\boldsymbol{Y}, U_Q) = \sum_{q=1}^{Q} f(\xi|\boldsymbol{Y}, M_q) P(M_q|\boldsymbol{Y}). \tag{3.1}$$

The posterior model probabilities in (3.1) are obtained through Bayes' formula:

$$P(M_q|\boldsymbol{Y}) = \frac{m(\boldsymbol{Y}|M_q)P(M_q)}{\sum_{q=1}^{Q} m(\boldsymbol{Y}|M_q)P(M_q)}, \tag{3.2}$$

where $P(M_q)$ are the prior model probabilities and $m(\boldsymbol{Y}|M_q)$ are the marginal likelihoods for each qth model. If the risk assessor has no preference for any particular model, a reasonable default views the M_qs as equally valid, so that $P(M_q) = 1/Q \quad \forall q$. The posterior model probabilities are then simply $P(M_q|\boldsymbol{Y}) = \frac{m(\boldsymbol{Y}|M_q)}{\sum_{q=1}^{Q} m(\boldsymbol{Y}|M_q)}$.

Clearly, the challenge with this sort of Bayesian model averaging (BMA) is to estimate the $P(M_q|\boldsymbol{Y})$ values. This requires accurate estimation of the marginal likelihood $m(\boldsymbol{Y}|M_q)$. Unfortunately, under our reparameterized hierarchy the associated integral is intractable. To approximate it, we employ the geometric bridge sampler of Meng and Wong (1996); also see Lopes and West (2004). Bridge sampling is relatively convenient and quite suitable for estimating marginal likelihoods from an AM sample. (Greater detail is provided in the Supplementary Document.) As pointed out by Hoeting *et al.* (1999), under a squared error loss function the BMA estimate of ξ is then straightforward to calculate. Simply take the weighted average of the individual-model, posterior sample mean estimates of $\xi_{100\text{BMR}}$ across the uncertainty class:

$$\bar{\xi}_{100\text{BMR}} = \sum_{q=1}^{Q} w_q \hat{\xi}_{100\text{BMR};q}\,, \tag{3.3}$$

where the weights are the posterior model probabilities, $w_q = P(M_q|\boldsymbol{Y})$. The primary focus in benchmark analysis is, however, the BMDL. To find a Bayesian BMA-based 95% BMDL we write $P(\xi \leq \underline{\xi}_{100\text{BMR}}|\boldsymbol{Y}, U_Q) = 0.05$, where ξ has p.d.f. defined by (3.1). From our AM sample, we approximate this via

$$P(\xi \leq \underline{\xi}_{100\text{BMR}}|\boldsymbol{Y}, U_Q) \approx \sum_{q=1}^{Q} w_q \frac{1}{K_q^*} \sum_{j=1}^{K_q^*} I\{\xi_j \leq \underline{\xi}_{100\text{BMR}}|M_q\},$$

where K_q^* is the size of the AM sample from model M_q (after burn-in) and I denotes the indicator function. Thus the estimated BMA BMDL satisfies the equation

$$\sum_{q=1}^{Q} w_q \frac{1}{K_q^*} \sum_{j=1}^{K_q^*} I\{\xi_j \leq \underline{\xi}_{100\text{BMR}}|M_q\} = 0.05.$$

The solution is found numerically. (Again, technical details are given in the Supplementary Document.)

4 Example: Benchmarking Carcinogenicity of Cumene

To illustrate our hierarchical BMA approach with the $Q = 8$ reparameterized models in §2.2, we returned to the cumene carcinogenicity data in Table 1. The BMR was set to the standard default level of 0.10 (U.S. EPA, 2012), and for the computations all experimental doses were scaled such that the maximum dose was equal to 1. With input from domain experts, we based prior elicitation for ξ and γ_0 on existing background in the toxicological literature. (Specifics are given in the Supplemental Document.) This led to the prior distributions $\xi \sim IG(0.53, 0.13)$ and $\gamma_0 \sim Beta(1.36, 12.31)$. For the three-parameter models, we required an additional prior specification for $\gamma_1 = R(d_m)$ at $d_m = 1$ (i.e., $d = 500$ ppm on the original scale). Unfortunately, no prior information was available on potential response of the tested animals at this (or any listed) dose of cumene (U.S. NTP, 2009). Thus we defaulted to use of an objective prior: $\gamma_1 \sim Beta(\frac{1}{2}, \frac{1}{2})$.

With this prior structure in place, we generated a Monte Carlo AM sample for each of the eight models. We encountered no algorithm failures, and the eight generated chains all passed our convergence diagnostic tests. This produced burn-ins of $10,000$ initial iterations for each chain, allowing us to operate with $90,000$ AM draws for each model. Using the methods described above, this led to model-specific, posterior-mean benchmark concentrations (BMCs) as reported in Table 3 (all final BMC values are rescaled to the original dose metric). The table also lists the corresponding 95% model-specific BMCLs, along with the posterior model probabilities/weights, $w_q = P(M_q|\boldsymbol{Y})$, assuming uniform prior model probabilities for each model. The eight BMC were then used to calculate the BMA BMC $\bar{\xi}_{10} = 27.4074$ ppm by using (3.3). The corresponding 95% BMCL is $\underline{\xi}_{10} = 15.1927$ ppm. These two values are also reported at the bottom of Table 3.

Table 3: BMC estimates based on posterior means and 95% BMCLs (in ppm) from each reparameterized model in §2.2, along with corresponding Bayesian model averaged (BMA) BMDL, for cumene carcinogenesis example. The BMR is set to 0.10.

Model	BMC_{10}	BMCL_{10}	$P(M_q\|\boldsymbol{Y})$
M_1	43.2752	35.5991	0.00044
M_2	44.7192	37.6845	0.00005
M_3	18.0881	14.7567	0.22887
M_4	76.3691	66.2304	0.00000
M_5	21.2154	16.2568	0.01356
M_6	31.1642	15.9229	0.36905
M_7	30.1092	15.3244	0.34956
M_8	24.2385	17.0606	0.03846
BMA	27.4074	15.1927	

From the table, we see that the posterior model probabilities vary widely with these data, suggesting that certain models may provide better-quality estimates than others. The three-parameter log-logistic and log-probit give the highest weights, both near 35%, followed by the

two-parameter quantal linear model at 23%. All other models show posterior probabilities below 5%. The BMA 95% BMCL is 15.1927 ppm, lying within the range of the individual model BMCLs and closer to the individual lower limits associated with higher-weight models.

The ramifications with these data for the risk analyst are substantial: had a choice of the poorly-fitting logistic, probit, or quantal-quadratic models been made for analyzing these data, the consequent BMC and BMCL would have been far too large. By integrating information across the various models, however, a more-tempered, model-robust estimate is produced from which further risk analytic calculations on cumene carcinogenicity can be conducted. Corroborating reports by many who have come before, we find that BMA adjustment frees the risk assessor from the selection biases, model inadequacies, and inferential uncertainties one encounters when committing to only a single parametric model to perform the benchmark analysis.

5 Performance Evaluations

5.1 Simulation design

To explore the features of our Bayesian model-averaged BMD/BMDLs in further detail, we conducted a short simulation study. We fixed the BMR at the standard level of BMR = 0.10 (U.S. EPA, 2012) and operated at a credible level of 95%. The doses were set to four levels: $d_1 = 0$, $d_2 = 0.25$, $d_3 = 0.5$, $d_4 = 1$, corresponding to a standard design in cancer risk experimentation (Portier, 1994). Equal numbers of subjects, $N_i = N$, were taken at each dose group. We considered three different per-dose sample sizes: $N = 25$, 50, or 1000; the latter approximates a 'large-sample' setting, while the former two are more commonly seen with toxicological investigations such as that in the cumene carcinogenicity example. As throughout, all of our calculations were performed within the `R` programming environment (R Development Core Team, 2012).

For the true dose-response model $R(d)$, we used each of the eight quantal-response functions in Table 2. We considered two different dose-response patterns, from a larger collection studied by West *et al.* (2012). The first pattern (P-I) was a moderately increasing response with $\gamma_0 = R(0) = 0.05$, $R(\frac{1}{2}) = 0.30$, and $R(1) = 0.50$. The second pattern (P-II) was more-broadly increasing, with $\gamma_0 = R(0) = 0.10$, $R(\frac{1}{2}) = 0.50$, and $R(1) = 0.90$. The resulting parameter configurations for the various models are given in Table 4.

At each of the 8 (models) $\times$ 2 (configurations) $\times$ 3 (sample sizes) = 48 combinations, we simulated 2000 pseudo-binomial quantal-response data sets via `R`'s `rbinom` function. Then for each data set, we generated eight AM samples, one from each model M_q. We then calculated the model-specific 95% BMDL $\underline{\xi}_{10;q}$ as the lower 5^{th} percentile of each model's AM sample. For simplicity, no prior elicitation was applied in the simulations and so we employed only objective priors for ξ, γ_0 and γ_1, i.e., $\xi \sim IG(0.001, 0.001)$, $\gamma_0 \sim Beta(\frac{1}{2}, \frac{1}{2})$, and $\gamma_1 \sim Beta(\frac{1}{2}, \frac{1}{2})$, to represent parameter uncertainty.

Assuming uniform prior model probabilities, the posterior model probabilities were calculated using geometric bridge sampling. We employed these posterior probabilities $P(M_q|\boldsymbol{Y})$ as weights and computed the BMA BMDLs according to the methods described in §3.

Table 4: Models and configurations for simulation study. Model codes are from Table 2

		Dose-Response Model							
Configuration	Parameters	M_1	M_2	M_3	M_4	M_5	M_6	M_7	M_8
P-I	γ_0	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05
	ξ	0.3974	0.3567	0.1642	0.4052	0.1783	0.2083	0.2267	0.1852
	γ_1	—	—	—	—	0.50	0.50	0.50	0.50
P-II	γ_0	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10
	ξ	0.1700	0.1575	0.0480	0.2190	0.1925	0.2760	0.2794	0.2025
	γ_1	—	—	—	—	0.90	0.90	0.90	0.90

5.2 Simulation Results

We studied how the BMA BMDL compared to the corresponding generating value of ξ_{10} in Table 4. In some sense, we expect 95% of these lower credible limits to lie below ξ_{10} and using our simulations as a guide, we queried how often this occurred. Figure 1 provides representative boxplots of the 2000 simulated BMA BMDLs under model M_6 (log-logistic) and configuration P-I at the popular per-dose sample size of $N = 50$. Mimicking a device employed by West *et al.* (2012), the boxplots are asymmetrically modified so that their upper whiskers stop at the 95th percentile of the 2000 simulated BMA BMDLs. (The lower whiskers rest at the minimum BMDL. The hinges and median bar are the usual quartiles.) Thus the 'goal' for each boxplot is to locate its upper whisker as close to, but not greatly exceeding, the generating value of ξ_{10}. A horizontal dashed line in the figure marks this ξ_{10} target.

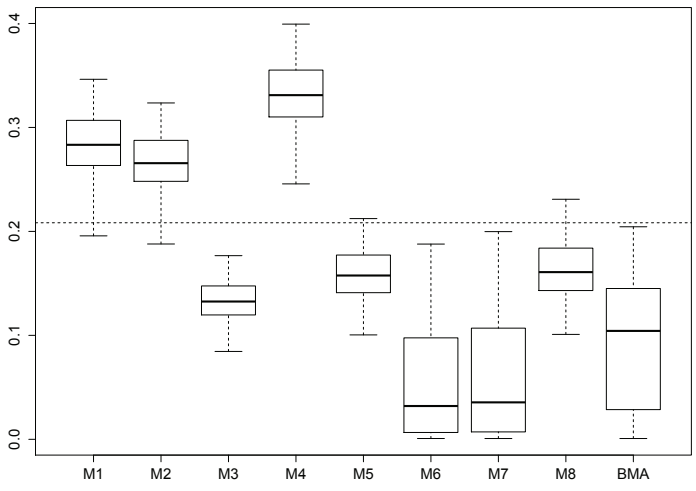

Figure 1: Modified Box plots for 95% individual-model BMDLs and BMA BMDL using simulated data from model M_6, configuration P-I, and sample size $N = 50$. (See text for details.) BMR is set to 0.10. Dashed horizontal line indicates target BMD_{10} under this model configuration.

The eight modified boxplots in Figure 1 correspond to model-specific 95% BMDLs calculated under each of the eight models in Table 2. The modified boxplot at far right gives the result for our 95% BMA BMDL. The graphic illustrates the consequences and ambiguities of single-model uncertainty when calculating Bayesian BMDLs. Recall that M_6 is the generating model and as anticipated, the M_6 boxplot for our single-model Bayesian BMDL displays acceptable characteristics: its modified upper whisker is slightly below the ξ_{10} target. It also exhibits a right skew, however, with median BMDL conservatively lower than any of the others in the figure.

The figure also illustrates the perplexing operating characteristics of single-model BMDLs under model misspecification: when M_1, M_2, and (particulary) M_4 are employed individually in the hierarchy, their BMDLs badly overestimate the ξ_{10} target. BMDLs using models M_5 and M_8 are somewhat more-stable; BMDLs using model M_3 show the least variation while locating conservatively below the ξ_{10} target. Analogous instances of stable/unstable single-model BMDLs occurred for all the model-configuration combinations we studied, but with no clear pattern as to which models operated well or poorly when misspecified. [Complete results from all 48 configurations are available in Fang (2014).]

On balance, however, the modified boxplot for the BMA BMDL at the far right of Figure 1 displays reasonable operating characteristics. As desired, its modified upper whisker lies just below the target ξ_{10}, and it locates a broader collection of $\underline{\xi}_{10}$ limits closer to that target, without exceeding it, than even the correct-model M_6 boxplot. (Admittedly, the M_3 and perhaps M_5 boxplots display even better performance here; however, these models did not exhibit consistently enhanced performance across all the configurations we studied.)

We find that unequivocal commitment to a specific dose-response model when there is any possibility of it being misspecified can lead, as often as not, to substantial overestimation of the benchmark point(s). Misspecifying the model sometimes produces acceptable lower limits, but we were not able to identify any predictable pattern of such among our simulation results. [This corroborates similar indications by West *et al.* (2012) for single-model, frequentist BMDL calculations.] By contrast, the Bayesian BMA BMDL provides a reasonable compromise.

Figure 2 extends this analysis to the large-sample setting, displaying a similar graphic when the per-dose sample size is $N = 1000$. Here, we highlight results for model M_3 under configuration P-II. As we might expect with such a large N, no single-model fit exhibits acceptable performance, save that of the correct model (M_3) and possibly models M_5 and M_8. (Notice that model M_3 is a special case of both M_5, with $\beta_2 = 0$, and M_8, with $\beta_1 = 1$, in Table 2. As in Figure 1, these three models sometimes perform similarly, although we also found cases where their operating characteristics diverged.) Moving to the BMA BMDL, however, overcomes these negative characteristics: the BMA BMDL boxplot at far right is almost identical to that for the correct-model M_3 boxplot. Here again, the BMA BMDL provides valuable robustness in the presence of model uncertainty/misspecification.

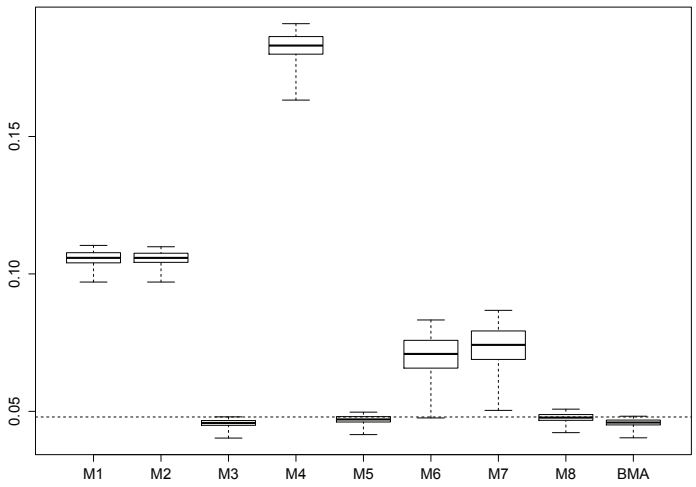

Figure 2: Modified Box plots for 95% individual-model BMDLs and BMA BMDL using simulated data from model M_3, configuration P-II, and sample size $N = 1000$. (See text for details.) BMR is set to 0.10. Dashed horizontal line indicates target BMD_{10} under this model configuration.

6 Discussion

Herein, we consider a strategy for model averaging within a hierarchical Bayesian framework when estimating benchmark doses (BMDs) in quantitative risk analysis. Placing emphasis on biomedical risk assessment, our approach estimates the BMD from a series of reparameterized quantal-response models with meaningful parameters (including the BMD itself) and accounts for potential model uncertainty via Bayesian model averaging (BMA). A mixture posterior density is constructed for the BMD, and is used to find the BMA point estimate (Hoeting *et al.*, 1999). The BMA lower credible limit (BMDL) is estimated as the lower αth percentile of the mixture posterior. Risk analysts can apply this method to avoid concerns of model uncertainty in multimodel problems, and to construct inferences on the BMD by incorporating prior knowledge for the uncertainty associated with pertinent model parameters.

Of course, some caveats and qualifications are in order. Our reparameterized dose-response models are highly complex and for the three-parameter forms they can become rather unwieldy. For instance, Fang (2014) notes that reparameterizations using γ_1 can produce unstable estimates when the highest dose level in the design is very close to the estimated BMD. In a series of additional simulations (results not shown) we constructed such a parameter setting where the generating value of ξ_{10} was slightly larger than d_m. We found that use of the three-parameter models often resulted in algorithm failures when an individual-model $\hat{\xi}$ drew near to d_m. Thus in cases where the BMD is felt to be near the highest dose—as determined, e.g., from elicited prior information—we recommend revising the reparameterization for γ_1 to define it as the response at some lower dose, say, $R(d_2)$.

This can ameliorate the instabilities that occur when $\hat{\xi}$ is near d_m.

It is also of interest to investigate how our approach operates under different design configurations. We focused on a geometric, four-dose design, arguably the quintessential standard in cancer and laboratory-animal toxicity testing. We may gain greater information about the pattern of dose response and therefore about the BMD, however, if we increase the number of doses and/or change the dose spacings. Experimental design for dose-response studies with focus on the BMD is an emerging area in the statistical literature (Muri *et al.*, 2009; Öberg, 2010; Sand *et al.*, 2008; Shao and Small, 2012) and how to optimally design/allocate experimental resources for BMD estimation and inferences under a Bayesian paradigm is an emerging question.

Acknowledgements

Thanks are due Drs. Katherine Y. Barnes, Anton Westveld and D. Dean Billheimer for their helpful suggestions during the preparation of this material. The results represent a portion of the first author's Ph.D. dissertation with the University of Arizona Graduate Interdisciplinary Program in Statistics.

References

Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In *Proceedings of the Second International Symposium on Information Theory* (Petrov B. N. and Csaki B., eds), 267–281. Akademiai Kiado, Budapest.

Andrieu, C. and Thoms, J. (2008). A tutorial on adaptive MCMC. *Statistics and Computing* **18**, 343–383.

Armitage, P. and Doll, R. (1954). The age distribution of cancer and a multi-stage theory of carcinogenesis. *British Journal of Cancer* **8**, 1–12.

Bailer, A. J., Noble, R. B. and Wheeler, M. W. (2005). Model uncertainty and risk estimation for experimental studies of quantal responses. *Risk Analysis* **25**, 291–299.

Casella, G. and Berger, R. L. (2002). *Statistical Inference*, 2nd edn. Pacific Grove, CA: Duxbury.

Christensen, R., Johnson, W. O., Branscum, A. J. and Hanson, T. E. (2011). *Bayesian Ideas and Data Analysis: An Introduction for Scientists and Statisticians.* Boca Raton, FL: Chapman & Hall/CRC Press.

Crump, K. S. (1984). A new method for determining allowable daily intake. *Fundamental and Applied Toxicology* **4**, 854–871.

Crump, K. S. (1995). Calculation of benchmark doses from continuous data. *Risk Analysis* **15**, 79–89.

Davis, J. A., Gift, J. S. and Zhao, Q. J. (2012). Introduction to benchmark dose methods and U.S. EPA's Benchmark Dose Software (BMDS) version 2.1.1. *Toxicology and Applied Pharmacology* **254**, 181–191.

European Union (2003). *Technical Guidance Document (TGD) on Risk Assessment of Chemical Substances following European Regulations and Directives, Parts I-IV*. Technical Report number EUR 20418 EN/1-4. Ispra, Italy: European Chemicals Bureau (ECB).

Fang, Q. (2014). *Hierarchical Bayesian Benchmark Risk Analysis*. Ph.D. thesis, Interdisciplinary Program in Statistics, University of Arizona, Tucson, AZ.

Fang, Q. and Piegorsch, W. W. (2014). Bayesian Benchmark Dose Analysis. Submitted.

Faustman, E. M. and Bartell, S. M. (1997). Review of noncancer risk assessment: Applications of benchmark dose methods. *Human and Ecological Risk Assessment* **3**, 893–920.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In *Bayesian Statistics* **4** (Bernardo, J. M., Berger, J. O., Dawid, A. P. and Smith, A. F. M., eds.), 169–193. Oxford University Press, Oxford.

Guha, N., Roy, A., Kopylev, L., Fox, J. Spassova, M. and White P. (2013). Nonparametric Bayesian Methods for Benchmark Dose Estimation. *Risk Analysis* **33**, 1608–1619.

Hoeting, J. A., Madigan, D., Raftery, A. E. and Volinsky, C. T. (1999). Bayesian model averaging: A tutorial. *Statistical Science* **14**, 382–401. (corr. **15**, 193–195).

Kang, S. H., Kodell. R. L. and Chen, J. J. (2000). Incorporating model uncertainties along with data uncertainties in microbial risk assessment. *Regulatory Toxicology and Pharmacology* **32**, 68–72.

Lambert, P. C., Sutton, A. J., Burton, P. R., Abrams, K. R. and Jones, D. R. (2005). How vague is vague? A simulation study of the impact of the use of vague prior distributions in MCMC using WinBUGS. *Statistics in Medicine* **24**, 2401–2428.

Lopes, H. F. and West, M. (2004). Bayesian model assessment in factor analysis. *Statistica Sinica* **14**, 41–67.

Meng, X. and Wong W. H. (1996). Simulating ratios of normalizing constants via a simple identity: A theoretical exploration. *Statistica Sinica* **6**, 831–860.

Morales, K. H., Ibrahim, J. G., Chen, C.-J. and Ryan, L. M. (2006). Bayesian model averaging with applications to benchmark dose estimation for arsenic in drinking water. *Journal of the American Statistical Association* **101**, 9–17.

Muri, S. D., Schlatter, J. R. and Brüschweiler, B. J. (2009). The benchmark dose approach in food risk assessment: Is it applicable and worthwhile? *Food and Chemical Toxicology* **47**, 2906–2925.

Naufal, Z., Kathman, S. and Wilson, C. (2009). Bayesian derivation of an oral cancer slope factor distribution for 4-(methylnitrosamino)-1-(3-pyridyl)-1-butanone (NNK). *Regulatory Toxicology and Pharmacology* **55**, 69–75.

Nitcheva, D. K., Piegorsch, W. W. and West, R. W. (2007). On use of the multistage dose-response model for assessing laboratory animal carcinogenicity. *Regulatory Toxicology and Pharmacology* **48**, 135–147.

Öberg M. (2010). Benchmark dose approaches in chemical health risk assessment in relation to number and distress of laboratory animals. *Regulatory Toxicology and Pharmacology* **58**, 451–454.

OECD (2006). *Current Approaches in the Statistical Analysis of Ecotoxicity Data: A Guidance to Application, Series on Testing and Assessment No. 54.* Paris: Environment Directorate, Organisation For Economic Co-Operation and Development.

OECD (2008). *Draft Guidance Document on the Performance of Chronic Toxicity and Carcinogenicity Studies, Supporting TG 451, 452 and 453.* Paris: Organisation For Economic Co-Operation and Development.

O'Hagan, A. (1994). *Kendall's Advanced Theory of Statistics, Volume 2B, Bayesian Inference*, 2nd edn. London: Edward Arnold.

Piegorsch, W. W. and Bailer, A. J. (2005). *Analyzing Environmental Data.* Chichester: John Wiley & Sons.

Piegorsch, W. W., An, L., Wickens, A. A., West, R. W. Peña, E. A. and Wu, W. (2013). Information-theoretic model-averaged benchmark dose analysis in environmental risk assessment. *Environmetrics*, **24**, 143–157.

Portier, C. J. (1994). Biostatistical issues in the design and analysis of animal carcinogenicity experiments. *Environmental Health Perspectives* **102**, Suppl. 1, 5–8.

R Development Core Team. (2012). *R: A Language and Environment for Statistical Computing.* Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0.

Robert, C. P. and Casella, G. (2011). A history of Markov chain Monte Carlo: subjective recollections from incomplete data. *Statistical Science* **26**, 102–115.

Sand, S., Victorin, K. and Falk Filipsson, A. (2008). The current state of knowledge on the use of the benchmark dose concept in risk assessment. *Journal of Applied Toxicology* **28**, 405–421.

Shao, K. (2012). A comparison of three methods for integrating historical information for Bayesian model averaged benchmark dose estimation. *Environmental Toxicology and Pharmacology* **34**, 288–296.

Shao, K. and Gift, J. S. (2014). Model uncertainty and Bayesian model averaged benchmark dose estimation for continuous data. *Risk Analysis* **34**, 101–120.

Shao, K. and Small, M. J. (2011). Potential uncertainty reduction in model-averaged benchmark dose estimates informed by an additional dose study. *Risk Analysis* **31**, 1561–1575.

Shao, K. and Small, M. J. (2012). Statistical evaluation of toxicological experimental design for Bayesian model averaged benchmark dose estimation with dichotomous data. *Human and Ecological Risk Assessment* **18**, 1096–1119.

Stern, A. H. (2008). Environmental health risk assessment. In *Encyclopedia of Quantitative Risk Analysis and Assessment* **2** (Melnick, E. L. and Everitt, B. S., eds.), 580–589. John Wiley & Sons, Chichester.

U.S. EPA (2012). *Benchmark Dose Technical Guidance Document*. Technical Report number EPA/100/R-12/001. Washington, DC: U.S. Environmental Protection Agency.

U.S. General Accounting Office (2001). *Chemical Risk Assessment. Selected Federal Agencies' Procedures, Assumptions, and Policies*. Report to Congressional Requesters number GAO-01-810. Washington, DC: U.S. General Accounting Office.

U.S. National Toxicology Program (2009). *Toxicology and Carcinogenesis Studies of Cumene (CAS NO. 98-82-8) in F344/N Rats and B6C3F$_1$ Mice*. Technical Report number 542. Research Triangle Park, NC: U.S. Department of Health and Human Services, Public Health Service.

West, R. W., Pigorsch, W. W., Peña, E. A., An, L., Wu, W., Wickens, A. A., Xiong, H., Chen, W. (2012). The impact of model uncertainty on benchmark dose estimation. *Environmetrics* **23**, 706–716.

Wheeler, M. W. and Bailer, A. J. (2009). Benchmark dose estimation incorporating multiple data sources. *Risk Analysis* **29**, 249–256.

Wheeler, M. W. and Bailer, A. J. (2012). Monotonic Bayesian semiparametric benchmark dose analysis. *Risk Analysis* **32**, 1207–1218.

arXiv:1402.3896v1 [stat.ME] 17 Feb 2014

Supplement to 'Bayesian Model-Averaged Benchmark Dose Analysis Via Reparameterized Quantal-Response Model'

Qijun Fang[a], **Walter W. Piegorsch**[a,b], **Susan J. Simmons**[c], **Xiaosong Li**[c], **Cuixian Chen**[c] and **Yishi Wang**[c]
[a]*Program in Statistics* and [b]*BIO5 Institute*
University of Arizona, Tucson, AZ, USA
[c]*Department of Mathematics and Statistics*
University of North Carolina at Wilmington, Wilmington, NC, USA

This supplementary document provides supporting material for *Bayesian Model-Averaged Benchmark Dose Analysis Via Reparameterized Quantal-Response Models* by Qijun Fang et. al., (the 'main document'). The various sections below address a variety of supplemental/supporting topics, and are not intended to flow naturally between each other. They are, however, presented in roughly the same order in which their counterpart topics appeared in the main document. As there, we denote the Benchmark Dose (BMD) target parameter as ξ at a Benchmark Response (BMR) between 0 and 1, the background response probability nuisance parameter as $\gamma_0 = R(0)$ and the response at the highest dose, d_m, as $\gamma_1 = R(d_m)$.

1 Example of reparameterizing Model M_5, the two-stage model

Here we study the two-stage model (M_5) to illustrate the reparameterization process, originally proposed in Fang and Piegorsch (2014). The risk function is

$$R(d) = 1 - \exp(-\beta_0 - \beta_1 d - \beta_2 d^2), \tag{S.1}$$

where $\beta_j \geq 0$ $(j = 0, 1, 2)$. With this, $\gamma_0 = R(0) = 1 - \exp(-\beta_0)$, hence $\exp(-\beta_0) = 1 - \gamma_0$. Substitute $\exp(-\beta_0) = 1 - \gamma_0$ into the risk function to find

$$R(d) = \gamma_0 + (1 - \gamma_0)[1 - \exp\{-\beta_1 d - \beta_2 d^2\}]. \tag{S.2}$$

Next, we reexpress β_1 and β_2 in terms of ξ, γ_0 and γ_1. Begin with γ_1, expressed in terms of γ_0, β_1 and β_2. From the definition of γ_1, we know

$$\gamma_1 = \gamma_0 + (1 - \gamma_0)[1 - \exp\{-\beta_1 d_m - \beta_2 d_m^2\}],$$

where d_m denotes the mth (highest) dose level. Let $\Gamma_5 = \log\left(\frac{1-\gamma_1}{1-\gamma_0}\right)$. After some algebra, we obtain

$$\beta_2 d_m^2 + \beta_1 d_m + \Gamma_5 = 0. \tag{S.3}$$

Now, Piegorsch and Bailer (2005, Ex. 4.12) show that the BMD for the two-stage model can be found to be

$$\xi = \frac{-\beta_1 + \sqrt{\beta_1^2 - 4\beta_2 \ln(1 - \text{BMR})}}{2\beta_2}. \tag{S.4}$$

For simplicity, denote $C_5 = -\ln(1 - \text{BMR})$. Then from Equation (S.4), we have

$$\beta_2 \xi^2 + \beta_1 \xi - C_5 = 0. \tag{S.5}$$

If we solve Equation (S.3) for β_1, we find $\beta_1 = (-\Gamma_5 - \beta_2 d_m^2)/d_m$. Substitute the result into (S.5) to achieve

$$\beta_2 = \frac{\Gamma_5 \xi + C_5 d_m}{\xi d_m (\xi - d_m)}. \tag{S.6}$$

Similarly, if we solve Equation (S.3) for β_2, then $\beta_2 = (-\Gamma_5 - \beta_1 d_m)/d_m^2$. Substitute the result into (S.5) to achieve

$$\beta_1 = \frac{C_5 d_m^2 + \Gamma_5 \xi^2}{\xi d_m (d_m - \xi)}. \tag{S.7}$$

So far, we have rewritten β_1 and β_2 in terms of Γ_5 (or γ_0 and γ_1) and ξ. The reparameterized risk function is obtained by substituting these expressions into Equation (S.2). This produces

$$R(d) = \gamma_0 + (1 - \gamma_0)\left[1 - \exp\left\{\frac{C_5 d_m d(d_m - d) + \Gamma_5 \xi d(\xi - d)}{\xi d_m (\xi - d_m)}\right\}\right], \tag{S.8}$$

where $\Gamma_5 = \log\left(\frac{1-\gamma_1}{1-\gamma_0}\right)$ and $C_5 = -\log(1 - \text{BMR})$. Similar sorts of operations may be applied to reparameterize any of the quantal-response models in main document Table 2 (Fang, 2014).

2 Finding prior parameters from elicited quantiles

We give here technical aspects on derivation of the prior parameters α, β, ψ, ω, κ and λ for the prior densities $\xi \sim IG(\alpha, \beta)$, $\gamma_0 \sim Beta(\psi, \omega)$ and $\gamma_1 \sim Beta(\kappa, \lambda)$ used in the main document's hierarchical model. The elicited lower quartiles and medians $Q_{1\xi}, Q_{2\xi}$ for ξ, $Q_{1\gamma_0}, Q_{2\gamma_0}$ for γ_0, and $Q_{1\gamma_1}, Q_{2\gamma_1}$ for γ_1, respectively, are assumed given from domain-expert input and/or previous literature/data.

Start with the elicitation for ξ: by definition, $Q_{1\xi}$ and $Q_{2\xi}$ satisfy

$$\int_0^{Q_{j\xi}} \frac{\beta^\alpha}{\Gamma(\alpha)} \xi^{(-\alpha-1)} \exp\left\{-\frac{\beta}{\xi}\right\} d\xi = \frac{j}{4},$$

$(j = 1, 2)$, where the integrand is the IG p.d.f. This establishes a system of non-linear equations for α and β. Unfortunately, no closed-form solution exists for the system, and so we turn to numerical methods. We employ a gradient method discussed by Barzilai and Borwein (1988) and implemented in the `R` statistical environment (R Development Core Team, 2012) via the `BB` package (Varadhan and Gilbert, 2009). This method employs iterative updating until half the L_2 norm of the system at the proposed solution is smaller than 10^{-10}. The iterative solution represents the 'elicited' values for α and β.

For γ_0, the elicited quartiles $Q_{1\gamma_0}$ and $Q_{2\gamma_0}$ satisfy

$$\int_0^{Q_{j\gamma_0}} \frac{\Gamma(\psi + \omega)}{\Gamma(\psi)\Gamma(\omega)} \gamma_0^{\psi-1} (1 - \gamma_0)^{\omega-1} d\gamma_0 = \frac{j}{4},$$

$(j = 1, 2)$, now applying the p.d.f. for the Beta prior. Here again, the resulting system of non-linear equations possesses no closed-form solution, so we appeal to the `BB` package in `R` to produce the elicited values for ϕ and ω. In similar fashion, and where necessary, given elicited quartiles $Q_{1\gamma_1}$ and $Q_{2\gamma_1}$ for γ_1, we solve for κ and λ in $\gamma_1 \sim Beta(\kappa, \lambda)$.

The resulting iterative solutions to these systems of equations produce values for the prior parameters α, β, ψ, ω, κ and λ which are then employed in our hierarchical model for the BMD.

3 Screen to check for 'data failures'

As described in Fang and Piegorsch (2014), some patterns of dose response can be problematic for benchmark dose estimation. Our experience shows that a shallow dose response can create

unstable frequentist $\hat{\xi}_{100\text{BMR}}$ estimates, and very shallow responses may cause the iterative model fit to fail entirely. Indeed, the EPA's BMDS software program will not estimate a BMD for a flat or negatively trending dose response (Wheeler and Bailer, 2009). For these reasons, we mark flat or decreasing-trend data as 'data failures' and not perform estimation for them.

To check the acceptability of the data before employing the algorithm, our simple screen is performed as follows: Over increasing doses $0 = d_1 < \ldots < d_i < \ldots < d_m$ $(i = 1, \ldots m)$, we first calculate the empirical extra risks:

$$\tilde{R}_E(d_i) = \frac{\frac{Y_i}{N_i} - \frac{Y_1}{N_1}}{1 - \frac{Y_1}{N_1}}.$$

We then connect each point $(d_i, \tilde{R}_E(d_i))$ to the origin point and find the largest slope, $S_{\max}$, among all $m-1$ rays. If $S_{\max} \leq 0$, no increasing trend is evidenced and we mark this as a 'data failure.' Notice that we do not perform a formal trend test to detect a decreasing trend (Wheeler and Bailer, 2009), but we do require that at least one estimated risk at some d_i $(i > 1)$ is higher than the estimated background risk.

4 The adaptive Metropolis (AM) algorithm

For data passing the simple 'data failure' screen in §3, we perform posterior simulation via a global AM procedure with componentwise adaptive scaling described by Andrieu and Thoms (2008). Their 'AM6' algorithm uses all previous iterations in the chain to update the variance-covariance matrix of the current bivariate proposal density. The operation is global in the sense that all the parameters (here, ξ and γ_0 for our two-parameter models and ξ, γ_0 and γ_1 for our three-parameter models) are updated simultaneously at each bivariate or trivariate draw. It is componentwise in the sense that the updating step sizes for ξ and γ_0 (or ξ, γ_0 and γ_1 for the three-parameter models) are controlled by separate adaptive scaling parameters. See Andrieu and Thoms (2008) for further details.

Denote $\boldsymbol{\theta}$ as the parameter vector $([\xi\ \gamma_0]^{\mathrm{T}}$ for the two-parameter case or $[\xi\ \gamma_0\ \gamma_1]^{\mathrm{T}}$ for the three-parameter case). Let $v_{k,1}, \ldots, v_{k,U}$ be adaptive scaling parameters and set $V_k = \text{diag}\,(v_{k,1}, \ldots, v_{k,U})$ at the kth draw in the Monte Carlo chain $(k = 1, \ldots, K)$. (U denotes the dimension for the parameter vector; $U = 2$ for our two-parameter models and $U = 3$ for the three-parameter models.) Take α as the acceptance probability in the normal random-walk Metropolis sampler, i.e.,

$$\alpha(\boldsymbol{\theta}_k, \boldsymbol{\theta}_k + Z_{k+1}) = \min\left\{1, \frac{\pi(\boldsymbol{\theta}_k|\boldsymbol{Y})}{\pi(\boldsymbol{\theta}_k + Z_{k+1}|\boldsymbol{Y})}\right\}.$$

For notation in what follows, set $\boldsymbol{e}_U$ as a length-U canonical vector with uth element equal to 1 and all others 0. Denote s_k as a step-size parameter that makes $\{s_k\}_{k=1}^{\infty}$ a decaying sequence; for example, the deterministic specification $s_k = 1/k^p$, with $p > 0$ is employed here. [Andrieu and Thoms (2008) recommend setting $p > 1/2$, while Vihola (2012) suggests $p = 2/3$. Here, Vihola's suggestion is followed and $s_k = k^{-2/3}$ is used.] A brief description of the AM algorithm follows.

- Find starting values for ξ, γ_0, and (where needed) γ_1 in the chain as follows: take the starting point for γ_0 by shrinking $\frac{Y_1}{N_1}$ away from 0 or 1 via $\frac{Y_1+0.25}{N_1+0.5}$. Similarly, take the starting point for γ_1 as $\frac{Y_j+0.25}{N_j+0.5}$, where the index j indicates the j^{th} data observation at which $S_{\max}$ is attained (see the 'simple screen' procedure described in §3). For the two-parameter models, the starting point for ξ is taken as $\text{BMR}/S_{\max}$. This is obtained by first treating $R_E(d) = S_{\max}d$ as the approximate extra risk function and then solving for d in the equation $\text{BMR} = S_{\max}d$. For

the three-parameter models, the starting point for ξ is taken as BMR/ $\left(\frac{\gamma_{10}-\gamma_{00}}{1-\gamma_{00}}\right)$ where γ_{00} and γ_{10} are the starting values for γ_0 and γ_1, respectively. These starting points collectively become the starting vector for the chain. Denote this as $\boldsymbol{\theta}_1$ in the algorithm and let $\mu_1 = \boldsymbol{\theta}_1$.

- Following Andrieu and Thoms (2008), let $\Sigma_1 = \mathbf{I}$ where $\mathbf{I}$ is the $U \times U$ identity matrix. Next, set $v_{1,1} = \ldots = v_{1,U} = \frac{2.38^2}{U}$ as the initial scaling parameters. Let $V_1 = \text{diag}\,(v_{1,1}, \ldots, v_{1,U})$, so that $V_1^{1/2}\mathbf{I}\Sigma_1^{1/2}$ is the initial variance-covariance matrix for use in the joint proposal distribution.

- Repeat for $k = 2, \ldots, K$:

  1. Given μ_{k-1}, Σ_{k-1} and $v_{k-1,1}, \ldots, v_{k-1,U}$, sample from a multivariate normal proposal distribution: $Z_k \sim N(0, V_{k-1}^{1/2}\Sigma_{k-1}V_{k-1}^{1/2})$ and set $\boldsymbol{\theta}_k = \boldsymbol{\theta}_{k-1} + Z_k$ with probability $\alpha(\boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_{k-1} + Z_k)$. Otherwise, set $\boldsymbol{\theta}_k = \boldsymbol{\theta}_{k-1}$.
  2. Update as follows for $u = 1, \ldots, U$:

$$\log(v_{k,u}) = \log(v_{k-1,u}) + s_k[\alpha(\boldsymbol{\theta}_{k-1}, \boldsymbol{\theta}_{k-1} + Z_k(u)e_u) - \bar{\alpha}_{**}],$$

     where $\bar{\alpha}_{**}$ is the acceptance rate of the candidate points generated at each iteration. Andrieu and Thoms (2008) suggest that $\bar{\alpha}_{**} = 0.44$ provides good performance; we followed their suggestion here.
  3. Update as follows for μ_k and Σ_k:

$$\mu_k = \mu_{k-1} + s_k(\boldsymbol{\theta}_k - \mu_{k-1}),$$

$$\Sigma_k = \Sigma_{k-1} + s_k[(\boldsymbol{\theta}_k - \mu_{k-1})(\boldsymbol{\theta}_k - \mu_{k-1})^T - \Sigma_{k-1}]$$

5 Monte Carlo 'burn-in' diagnostics

To account for potential instability in the early portions of the length-K, bivariate, AM chain from §4, we include a 'burn-in' for the Monte Carlo sample (Gelman *et al.*, 2004, §11.6). We couple this with the larger question of how to assess the chain's convergence. Many approaches exist for diagnosing MCMC convergence (Cowles and Carlin, 1996), and indeed, the issue is a topic of ongoing debate (Robert and Casella, 2004, §12.2). We mimic a method due to Geweke (1992), where early portions of the chain are sub-sampled and compared against latter portions to determine where the larger sample of K draws begins to exhibit stable performance. To approximate independence between the two sub-samples, we bifurcate the chain by a gap of no less than $\frac{K}{5}$ consecutive draws.

For summary diagnostics, we calculate the difference in arithmetic means between the two bifurcated sub-samples, and divide each difference by its standard error to produce a Z-statistic. The approximate standard error is taken as the square root of the sum of the variances of each mean. This is done separately for the ξ_k, γ_{0k}, and γ_{1k} components. To adjust for possible autocorrelation within each sub-sample, the individual variances are based on estimated spectral densities at frequency zero. Two parametric methods of estimating these spectral densities are considered here. The first method is discussed by Heidelberger and Welch (1981), which consists of fitting a generalized linear model to the periodogram of the batched time series. We implement this using `R`, via the `coda` package's `spectrum0` function. The second method begins by fitting the periodogram with an autogressive model, and then use the spectral density of the fitted AR model to approximate the spectral density at frequency zero. We again implement this method using `R`, now with

the `coda` package's `spectrum0.ar` function. These two methods produced similar estimates of the spectral density at zero. The Heidelberger and Welch method is fast and is the default method in the `coda` package, hence we also default to it first. The fit may fail due to high autocorrelation of the chain, however, and when this happens we then switch to the AR method for estimating the spectral density at zero.

To monitor if the pattern of association between any two parameters within the larger chain also exhibits reasonable stability, we include comparison of up to three pairs of covariances between two sub-samples (for the two-parameter models, only the covariance between ξ and γ_0 is considered). As a diagnostic measure we use the sample covariances from each sub-sample, which are

$$\widehat{\text{Cov}}(\xi, \gamma_0) = \frac{1}{L}\sum_{k=1}^{L}(\xi_k - \bar{\xi})(\gamma_{0k} - \bar{\gamma}_0),$$

$$\widehat{\text{Cov}}(\xi, \gamma_1) = \frac{1}{L}\sum_{k=1}^{L}(\xi_k - \bar{\xi})(\gamma_{1k} - \bar{\gamma}_1),$$

$$\widehat{\text{Cov}}(\gamma_0, \gamma_1) = \frac{1}{L}\sum_{k=1}^{L}(\gamma_{0k} - \bar{\gamma}_0)(\gamma_{1k} - \bar{\gamma}_1),$$

where $\bar{\xi}$, $\bar{\gamma}_0$ and $\bar{\gamma}_1$ are the pertinent arithmetic means within a bifurcated sub-sample of length L (see below). To find the approximate variance of each estimated covariance we calculate the quantities

$$\tau_k(\xi, \gamma_0) = (\xi_k - \bar{\xi})(\gamma_{0k} - \bar{\gamma}_0),$$
$$\tau_k(\xi, \gamma_1) = (\xi_k - \bar{\xi})(\gamma_{0k} - \bar{\gamma}_1),$$
$$\tau_k(\gamma_0, \gamma_1) = (\gamma_{0k} - \bar{\gamma}_0)(\gamma_{1k} - \bar{\gamma}_1),$$

across all values of the index k in the given sub-sample and then estimate the variance of $\tau_k(\xi, \gamma_0)$, $\tau_k(\gamma_0, \gamma_1)$ and $\tau_k(\gamma_0, \gamma_1)$ based on their estimated spectral densities at frequency zero.

Each such comparison is conducted over three individual, separated bifurcations of the full K-length sample: (i) the first 10% of the chain ($L = K/10$) vs. the final 50% ($L = K/2$), (ii) the first 20% of the chain ($L = K/5$) vs. the final 50%, and (iii) the first 30% of the chain ($L = 0.3K$) vs. the final 50%. We consider an individual diagnostic comparison as a 'pass' if the corresponding Z-statistic is less than 1.96 in absolute value. For the two-parameter models, to 'pass' the full diagnostic at each bifurcation, all three measures—mean of ξ, mean of γ_0, and covariance of $\{\xi,\gamma_0\}$—must individually pass; for the three-parameter models, in addition, the mean of γ_1, and the covariances of $\{\xi, \gamma_1\}$ and $\{\gamma_0, \gamma_1\}$ must also pass. The comparisons are performed in sequential order: if the 10%-vs.-50% diagnostic fails, then the 20%-vs.-50% diagnostic is conducted. If this fails, we move to the 30%-vs.-50% diagnostic. When a diagnostic passes, we view the indicated early portion of the chain as a reasonable burn-in. For notation, we let $K_0 < K$ be the resulting index that begins the retained portion of the chain. Thus, e.g., if the 10%-vs.-50% diagnostic fails but then the 20%-vs.-50% diagnostic passes, we take the first 20% of the chain as burn-in and use the remaining 80% of the chain as our Monte Carlo sample from $\pi(\xi, \gamma_0|\boldsymbol{Y})$ or $\pi(\xi, \gamma_0, \gamma_1|\boldsymbol{Y})$. If K =100,000, as used throughout, this gives K_0 =20,001.

If none of the three sequential diagnostics pass, we re-set the random-number generator's seed and reassess a new set of K Monte Carlo draws. If after five such re-starts the diagnostic continues to fail, we report an *algorithm failure*, and terminate the analysis with the given data set.

6 Bridge sampling for estimating marginal distributions

As a necessary step to perform Bayesian model averaging, we must calculate an approximation to marginal likelihoods of the form (with the two-parameter models)

$$m(\boldsymbol{Y}) = \int_0^\infty \int_0^1 f(\boldsymbol{Y}|\xi, \gamma_0)\pi(\xi|\alpha, \beta)\pi(\gamma_0|\psi, \omega)d\gamma_0 d\xi$$

for elicited or objective priors $\pi(\xi|\alpha, \beta)$ and/or $\pi(\gamma_0|\psi, \omega)$. To do so, we constructed a *geometric estimator* from our AM sample, appealing to the 'bridge sampling' method recommended by Meng and Wong (1996) and Lopes and West (2004). To approximate the marginal likelihood, an approximation $g(\cdot)$ to the joint posterior density of (ξ, γ_0) is first chosen. For simplicity, we took $g(\xi, \gamma_0)$ as the bivariate normal density function with mean set to the empirical mean vector and variance set to the empirical variance-covariance matrix of ξ and γ_0, estimated from the AM sample. Next, a set of $K^* = K - K_0$ bivariate vectors $\left\{[\xi_j^* \ \gamma_{0j}^*]^T\right\}_{j=1}^{K^*}$ are drawn from $g(\xi, \gamma_0)$. The marginal p.d.f. is then approximated by substituting the (retained) AM sample, $\{\xi_k, \gamma_{0k}\}_{k=K_0}^{K}$, and the generated sample from $g(\xi, \gamma_0)$, $\left\{[\xi_j^* \ \gamma_{0j}^*]^T\right\}_{j=1}^{K^*}$, into the ratio

$$\widehat{m}(\boldsymbol{Y}) = \frac{\frac{1}{K^*}\sum_{j=1}^{K^*}\left\{\pi\left(\xi_j^*|\alpha,\beta\right)\pi\left(\gamma_{0j}^*|\psi,\omega\right)f\left(\boldsymbol{Y}|\xi_j^*,\gamma_{0j}^*\right)/g\left(\xi_j^*,\gamma_{0j}^*\right)\right\}^{1/2}}{\frac{1}{K^*}\sum_{k=K_0}^{K}\left\{g\left(\xi_k,\gamma_{0k}\right)/\left[\pi\left(\xi_k|\alpha,\beta\right)\pi\left(\gamma_{0k}|\psi,\omega\right)f\left(\boldsymbol{Y}|\xi_k,\gamma_{0k}\right)\right]\right\}^{1/2}},$$

where $f(\cdot)$ denotes the binomial likelihood function and $\pi(\cdot)$ denotes the pertinent prior density. The above expression gives the approximated marginal likelihood for our two-parameter models. Marginal likelihoods for our three-parameter models are approximated in the same fashion by including the $\pi(\gamma_1|\kappa, \lambda)$ prior and using a trivariate approximating normal density $g(\xi, \gamma_0, \gamma_1)$.

7 BMA BMDL calculation

In the main document, we claimed that

$$P(\xi \leq \underline{\xi}_{100\text{BMR}}|\boldsymbol{Y}, U_Q) \approx \sum_{q=1}^{Q} w_q \frac{1}{K_q^*}\sum_{j=1}^{K_q^*} I\{\xi_j \leq \underline{\xi}_{100\text{BMR}}|M_q\}.$$

Here we give a simple proof.

Proof. Equation (3.1) in the main document defines the model-averaged posterior density for the BMD, $f(\xi|\boldsymbol{Y}, U_Q)$, as a mixture of marginal posterior densities for ξ under each qth model, $f(\xi|\boldsymbol{Y}, M_q)$ using each model's posterior probability $P(M_q|\boldsymbol{Y})$ as the corresponding weight, w_q.

Denote $F(\xi|\boldsymbol{Y}, U_Q)$ as the model-averaged posterior distribution function for the BMD and $F(\xi|\boldsymbol{Y}, M_q)$ as the marginal posterior distribution function for ξ under each model. By integrating both sides of Equation (3.1), we have

$$F(\xi|\boldsymbol{Y}, U_Q) = \sum_{q=1}^{Q} w_q F(\xi|\boldsymbol{Y}, M_q).$$

Now, substitute $\underline{\xi}_{100\text{BMR}}$ into the above equation to find

$$\begin{aligned} P(\xi \leq \underline{\xi}_{100\text{BMR}}|\boldsymbol{Y}, U_Q) &= F(\underline{\xi}_{100\text{BMR}}|\boldsymbol{Y}, U_Q) \\ &= \sum_{q=1}^{Q} w_q F(\underline{\xi}_{100\text{BMR}}|\boldsymbol{Y}, M_q) \\ &= \sum_{q=1}^{Q} w_q P(\xi \leq \underline{\xi}_{100\text{BMR}}|\boldsymbol{Y}, M_q). \end{aligned}$$

Clearly then, the probability, $P(\xi \leq \underline{\xi}_{100\text{BMR}}|\boldsymbol{Y}, M_q)$ can be approximated by the sample proportion $\frac{1}{K_q^*}\sum_{j=1}^{K_q^*} I\{\xi_j \leq \underline{\xi}_{100\text{BMR}}|M_q\}$, where K_q^* is the size of the AM sample under M_q (excluding burn-in) and I denotes the indication function. □

The BMA BMDL is then calculated by finding the root of equation

$$\frac{1}{K_q^*}\sum_{j=1}^{K_q^*} I\{\xi_j \leq \underline{\xi}_{100\text{BMR}}|M_q\} = 0.05.$$

To solve this equation numerically, we use the `R` `uniroot` function. This employs iterative updating until either an exact root is found or the change in root for one step of the algorithm is less than 2^{-13}.

8 Prior elicitation with the cumene data

To elicit the prior parameters for the cumene data in the main document, we applied input from domain experts and based the prior elicitation on existing background from the toxicological literature. We began with the BMD target parameter, ξ: an oral No Observed Adverse Effect Level (NOAEL) for long-term cumene exposure in rodents was given by the EPA IRIS website (`http://www.epa.gov/iris/subst/0306.htm`) as 110 mg/kg-day. Guidance for converting the oral mg/kg-day metric to inhalation ppm (as used in the original data Table from the main document) was available via conversion data given by the California OEHHA (2004): for rodents, 1 mg/kg-day works out to roughly 2.25 ppm. Arguing that a NOAEL is loosely equivalent to a BMD—although, see (Kodell, 2005) for an in-depth discussion—the median prior estimate was $Q_{2\xi}$ = 110 mg/kg-day = 247.5 ppm ≈ 250 ppm. For the first quartile, the EPA IRIS site gave a cumene inhalation NOAEL for rodents as 435 mg/cu.m, based on a shorter, 13-week study. The shorter exposure period was suited to provide only a conservative estimate on the BMD's central tendency; we translated this to use of the 435 mg/cu.m 13-week NOAEL as a prior estimate of the first quartile. For conversion to inhalation ppm, the EPA indicated that roughly 1 mg/cu.m = 0.204 ppm, so we took $Q_{1\xi}$ = 435 mg/cu.m = 88.75 ppm ≈ 90 ppm. We then divided the ppm dose by 500 ppm to standardize the ξ-scale to $0 \leq \xi \leq 1$. This produced prior first quartile and median estimates for $\pi(\xi|\alpha, \beta)$ as 0.18 and 0.5, respectively.

For γ_0, historical control data on alveolar/bronchiolar tumors in Female B6C3F$_1$ mice were given by U.S. NTP (2009, Table D3a). These provide direct prior information on the background risk. The historical data gave a median background response of $Q_{2\gamma_0} = 0.08$ and a lower quartile of $Q_{1\gamma_0} = 0.04$.

As noted in the main document, no elicitation for γ_1 was available. We found no indication of carcinogenic effects of cumene evidenced in laboratory mice prior to the NTP's toxicology study (U.S. NTP, 2009). As a result, we defaulted to the objective prior $\gamma_1 \sim Beta\left(\frac{1}{2}, \frac{1}{2}\right)$.

These quartile-based elicitations were then used to construct the prior densities $\xi \sim IG(\alpha, \beta)$ and $\gamma_0 \sim Beta(\psi, \omega)$ as described in §1, above.

References

Andrieu, C. and Thoms, J. (2008). A tutorial on adaptive MCMC. *Statistics and Computing* **18**, 343–383.

Barzilai, J. and Borwein, J. M. (1988). Two-point step size gradient methods. *IMA Journal of Numerical Analysis* **8**, 141–148.

California Office of Environmental Health Hazard Assessment (2004). *Proposition 65 Maximum Allowable Dose Level (MADL) for Reproductive Toxicity for Methyl Bromide as a Structural Fumigant*. Sacramento, CA: CA Office of Environmental Health Hazard Assessment (OEHHA), Reproductive and Cancer Hazard Assessment Section.

Cowles, M. K. and Carlin, B. P. (1996). Markov chain Monte Carlo convergence diagnostics: A comprarative review. *Journal of the American Statistical Association*, **91**, 883–904.

Fang, Q. (2014). *Hierarchical Bayesian Benchmark Risk Analysis*. Ph.D. thesis, Interdisciplinary Program in Statistics, University of Arizona, Tucson, AZ.

Fang, Q. and Piegorsch, W. W. (2014). Bayesian Benchmark Dose Analysis. Submitted.

Gelman, A., Carlin, B. J., Stein, H. S. and Rubin, D. B. (2004). *Bayesian Data Analysis*, 2nd edn. Boca Raton, FL: Chapman & Hall/CRC.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In Bernardo, J. M., Berger, J. O., Dawid, A. P. and Smith, A. F. M. (eds.), *Bayesian Statistics* **4**, 169–193. Oxford: Oxford University Press.

Heidelberger, P and Welch, P. D. (1981). A spectral method for confidence interval generation and run length control in simulations. *Communications of the ACM* **24**, 233–245.

Kodell, R. L. (2005). Managing uncertainty in health risk assessment. *International Journal of Risk Assessment and Management* **14**, 193–205.

Lopes, H. F. and West, M. (2004). Bayesian model assessment in factor analysis. *Statistica Sinica* **14**, 41–67.

Meng, X. and Wong W. H. (1996). Simulating ratios of normalizing constants via a simple identity: A theoretical exploration. *Statistica Sinica* **6**, 831–860.

Piegorsch, W. W. and Bailer, A. J. (2005). *Analyzing Environmental Data*. Chichester: John Wiley & Sons.

R Development Core Team. (2012). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0.

Robert, C. P. and Casella, G. (2004). *Monte Carlo Statistical Methods*, 2nd Edn. New York: Springer-Verlag.

U.S. National Toxicology Program (2009). *Toxicology and Carcinogenesis Studies of Cumene (CAS NO. 98-82-8) in F344/N Rats and B6C3F$_1$ Mice*. Technical Report number 542. Research Triangle Park, NC: U.S. Department of Health and Human Services, Public Health Service.

Varadhan, R. and Gilbert, P. (2009). BB: An R package for solving a large system of nonlinear equations and for optimizing a high-dimensional nonlinear objective function. *Journal of Statistical Software* **32**, 1–26.

Vihola, M. (2012) Robust adaptive Metropolis algorithm with coerced acceptance rate. *Statistics and Computing* **22**, 997–1008.

Wheeler, M. W. and Bailer, A. J. (2009). Comparing model averaging with other model selection strategies for benchmark dose estimation. *Environmental and Ecological Statistics* **16**, 37–51.